Supplementary Notes: **Details on the Supplementary Data 1, 3 and 4, on the Supplementary Figures 2, and on the Supplementary Program 1**

Supplementary Data 1: Locations (UTM zone 17 coordinates, elevation in meters), estimated mean annual dry season duration (days), mean annual precipitation (mm), and the number of complete years' data used (only data since 1960) for the 44 rainfall collection stations operated by the Panama Canal Authority[24].

Supplementary Data 3,4, Supplementary Figures 3: 44 of the experimental species were analyzed for their occurrence across the isthmus. Two species (*Piper trigonum* and *Anaxogorea panamensis*) did not occur in any of the 122 inventory sites, and two species (*Inga multijuga* and *Inga cocleensis*) were excluded because of uncertainties with species identification.

Supplementary Data 4: Species habitat response. The occurrence probabilities are k_{145} and k_{110}, estimated from the median kernel from the Gibbs sampler; the LogRatio is the log base 10 of that ratio (since the ratio k_{145}/k_{110} is the median of the ratios from the Gibbs sampler, it is not exactly the ratio of the displayed medians). Upper and lower ratios are the 2.5^{th} and 97.5^{th} percentiles of the ratio from the Gibbs chain. The bandwidths are the standard deviation of the Gaussian kernel, and are median plus 2.5^{th} and 97.5^{th} percentiles from the Gibbs chain. Low bandwidths indicate strong variation in species occurrence with drought; median and upper values > 100 indicate no drought response.

Supplementary Figures 3: Species occurrence fitted against dry season length. Vertical blue bars indicate inventories where the species was observed, red bars indicate inventories where the species was not observed. The black line is the prediction based on the median bandwidth (Supplementary Data 4); the 20 gray curves are drawn from bandwidths randomly chosen from the Gibbs sampler and illustrate statistical confidence in predicted occurrence. Vertical lines at $d = 110$ days and $d = 145$ days show occurrence probability toward the wet and dry ends of the gradient. In species where some of the lines are flat, a naive model of no drought response is statistically acceptable.

Supplementary Program 1: R scripts for running the Gibbs sampler to fit an optimized Gaussian kernel to species occurrence as a function of dry season duration. The main function is *fit.kernel.Gibbs.predictors* and will finish the entire process for a single species, given a vector of *responses* (here, 0s and 1s to indicate presence and absence of the species at 122 sites), a vector of *predictors* (here, dry season duration at the same 122 locations), and *start* (a starting value of the band width, set to 10 in our runs).